arXiv:1411.1211v1 [math.OC] 5 Nov 2014

Generic uniqueness of the bias vector of mean payoff zero-sum games

Marianne Akian, Stéphane Gaubert and Antoine Hochart

***Abstract*— Zero-sum mean payoff games can be studied by means of a nonlinear spectral problem. When the state space is finite, the latter consists in finding an eigenpair (u, λ) solution of $T(u) = \lambda \mathbf{1} + u$ where $T : \mathbb{R}^n \to \mathbb{R}^n$ is the Shapley (dynamic programming) operator, λ is a scalar, $\mathbf{1}$ is the unit vector, and $u \in \mathbb{R}^n$. The scalar λ yields the mean payoff per time unit, and the vector u, called the *bias*, allows one to determine optimal stationary strategies. The existence of the eigenpair (u, λ) is generally related to ergodicity conditions. A basic issue is to understand for which classes of games the bias vector is unique (up to an additive constant). In this paper, we consider perfect information zero-sum stochastic games with finite state and action spaces, thinking of the transition payments as variable parameters, transition probabilities being fixed. We identify structural conditions on the support of the transition probabilities which guarantee that the spectral problem is solvable for all values of the transition payments. Then, we show that the bias vector, thought of as a function of the transition payments, is generically unique (up to an additive constant). The proof uses techniques of max-plus (tropical) algebra and nonlinear Perron-Frobenius theory.**

I. Introduction

A. The mean payoff of a repeated game

Zero-sum repeated games describe long term interactions between two agents (called players) with opposite interests. We consider here perfect information games, in which each of the players chooses alternatively an action, being informed of the current state of the game and of the previous actions of the other player. These choices determine an instantaneous payment, as well as the next state, either by a deterministic or a stochastic process. We refer the reader to [1] for background on repeated games and stochastic games.

A zero-sum repeated stochastic game with perfect information is composed of
– a state space S.
– an action space $A_i \subset A$ for player MIN, depending on $i \in S$, and included in a given set A; at each stage, player MIN chooses an action $a \in A_i$ knowing the current state i and the past states and actions of both players.
– an action space $B_{i,a} \subset B$ for player MAX, depending on $i \in S$ and on $a \in A_i$, and included in a given set B; at each stage, player MAX chooses an action $b \in B_{i,a}$, knowing the current state i, the last action $a \in A_i$ chosen by player MIN and the past states and actions of both players.

M. Akian, S. Gaubert and A. Hochart are with INRIA Saclay-Ile-de-France and CMAP, Ecole Polytechnique, CNRS, CMAP, Route de Saclay, 91128 Palaiseau Cedex, France.

The last author is supported by a PhD fellowship of Fondation Mathématique Jacques Hadamard.

marianne.akian@inria.fr (M. Akian)
stephane.gaubert@inria.fr (S. Gaubert)
hochart@cmap.polytechnique.fr (A. Hochart)

– a transition payment $r_i^{ab} \in \mathbb{R}$ paid by player MIN to player MAX at each stage, given the current state $i \in S$ and the last actions of each of the players $a \in A_i$ and $b \in B_{i,a}$.
– a transition probability $P_i^{ab} \in \Delta(S)$, depending on the same data, where $\Delta(S)$ is the set of probabilities over S; when in state $i \in S$, if the actions a, b have been played, the next state is chosen according to the probability P_i^{ab}.

The game is said to be *deterministic* if each transition probability P_i^{ab} has only one state in its support, meaning that for every choice of $i \in S$, $a \in A_i$, and $b \in B_{i,a}$, we have $P_i(\{j\}) = 1$ for precisely one state $j \in S$. Given the above parameters, and a strategy of each player (that is a selection rule for the sequence of actions at each step of the game), the payoff of the game is an additive function of the transition payments, that player MIN intends to minimize and player MAX intends to maximize. In particular, the game with a fixed finite horizon k consists of k successive alternated moves of the players MIN and MAX, with a payoff

$$J_i^k = \mathbb{E}\left[\sum_{\ell=0}^{k-1} r_{\xi_\ell}^{\alpha_\ell \beta_\ell}\right], \tag{1}$$

where $\mathbb{E}$ denotes the expectation for the probability law of the process $(\xi_\ell, \alpha_\ell, \beta_\ell)_{\ell \geq 0}$ of state and actions of players MIN and MAX respectively, determined by the above transition probabilities and the strategies, given the initial state i.

Throughout the paper, we make the following assumptions:

Assumption A (finite state and action spaces):

i) The state space S is finite, say $S = \{1, \dots, n\}$.
ii) All the action spaces A_i and $B_{i,a}$ are nonempty finite sets.

The analysis of the above game involves the *Shapley* operator $T : \mathbb{R}^n \to \mathbb{R}^n$, which is such that, for all $x = (x_i)_{i \in S}$,

$$[T(x)]_i = \min_{a \in A_i} \max_{b \in B_{i,a}} \left(r_i^{ab} + P_i^{ab} x\right), \tag{2}$$

for all states $i \in S$. Note that Assumption A implies that the min and max in (2) are always attained. Also, the elements P of $\Delta(S)$ are seen as row vectors, $P = (P_j)_{j \in S}$, so that Px means $\sum_{j=1}^n P_j x_j$.

Given an initial state $i \in S$, the game in horizon k is known to have a value, denoted by $v_i^k \in \mathbb{R}$, which is equal to the minimum over all strategies of MIN of the maximum over all strategies of MAX of the payoff J_i^k of (1). In particular, the value vector $v^k = (v_i^k)_{1 \leq i \leq n} \in \mathbb{R}^n$ satisfies the dynamic programming recursion

$$v^0 = 0 \text{ and } v^{k+1} = T(v^k). \tag{3}$$

Note that, since all the action spaces are finite, the above game can be transformed into a game with same behavior and operator (2), but with action spaces independent of the state and the selected action (by duplicating some transition payments and probabilities). However, this would change the space of transition payments hence most of our results.

We will study here the asymptotics of the value vector as the horizon k of the game tends to infinity. In particular, we are interested in the following quantity, called the *mean payoff* vector (per time unit):

$$\chi(T) := \lim_{k\to\infty} v^k/k = \lim_{k\to\infty} T^k(0)/k \ . \tag{4}$$

The mean payoff $\chi(T)$ does exist when the action spaces are finite. Indeed, Kohlberg [2] has shown that T has then an invariant half-line, that is, there exist two vectors $x, \nu \in \mathbb{R}^n$ such that $T(x + t\nu) = x + (t + 1)\nu$ for every scalar t large enough. It can be checked that $\chi(T) = \nu$. Note that the existence of the mean payoff vector is also guaranteed more generally when T is semi-algebraic, see Neyman [3]. Other conditions, in terms of the operator $(\alpha, x) \mapsto \alpha T(\frac{1-\alpha}{\alpha}x),\ \ \alpha > 0$, have been given by Rosenberg and Sorin [4].

It is convenient to note that the limit $\lim_{k\to\infty} T^k(x)/k$ is independent of the choice of $x \in \mathbb{R}^n$. This follows readily from the elementary fact that T is sup-norm nonexpansive.

B. Nonlinear spectral problem

The existence of the mean payoff vector is guaranteed if we can find a vector $u \in \mathbb{R}^n$ and a scalar $\lambda \in \mathbb{R}$ solution of the following nonlinear spectral problem,

$$T(u) = \lambda\mathbf{1} + u, \tag{5}$$

where $\mathbf{1}$ is the unit vector of $\mathbb{R}^n$.

Indeed, in such a situation, we have $T^k(u) = k\lambda + u$, for all $k \geq 0$, and so, $\chi(T) = \lambda\mathbf{1}$, meaning that the mean payoff starting from state i, $[\chi(T)]_i$, is independent of the initial state i, and equal to λ. The scalar λ will be called the *eigenvalue* of T. It follows that the eigenvalue, if it exists, is necessarily unique (for $\chi(T)$ is uniquely defined). The vector u, called *bias vector* or *eigenvector*, gives optimal stationary strategies.

Since any Shapley operator (2) is additively homogeneous (that is, commutes with the addition of a constant), a bias vector is always defined up to an additive constant and it is interesting to understand when such a vector is unique (up to an additive constant).

For one player problems, i.e., for discrete optimal control, the nonlinear eigenproblem (5) (also known as the 'average case optimality equation') has been much studied, either in the deterministic or in the stochastic case (Markov decision problems). Then, the representation of bias vectors and their relation with optimal strategies is well understood.

For the deterministic case, the analysis relies on max-plus spectral theory, which goes back to the work of Romanovski [5], Gondran and Minoux [6] and Cuninghame-Green [7]. In [8], [9], [10] can be found more background on max-plus spectral theory. Kontorer and Yakovenko [11], and Kolokoltsov and Maslov [12], deal specially with infinite horizon optimization and mean payoff problems. The set of bias vectors has the structure of a max-plus (tropical) cone, i.e., it is invariant by max-plus linear combinations, and it has a unique minimal generating family consisting of certain 'extreme' generators, which can be identified by looking at the support of the maximizing measures in the linear programming formulation of the optimal control problem, or at the "recurrence points" of infinite optimal trajectories. We refer the reader to [13] and to the references therein for more information on max-plus spectral theory. A combinatorial interpretation of some of these results, in terms of polyhedral fans, has been recently given by Sturmfels and Tran [14]. The eigenproblem has also been studied for an infinite dimensional state space in the context of infinite dimensional max-plus spectral theory, see Akian, Gaubert and Walsh [15], and also in the setting of weak KAM theory, for which we refer the reader to the book of Fathi [16]. In the stochastic case, the structure of the set of bias vectors is still known, at least when the state space is finite, see Akian and Gaubert [17].

In the two player case, the structure of the set of bias vectors $u \in \mathbb{R}^n$ is less well known, although the description of this set remains a fundamental issue. In particular, the uniqueness of the bias vector (up to an additive constant) is an important matter for algorithmic purposes. Indeed, the nonuniqueness of the bias typically leads to numerical unstabilities or degeneracies. In particular, the standard Hoffman and Karp policy iteration algorithm [18] may fail to converge in situations in which the bias vector is not unique. Some refinements of the Hoffman and Karp scheme have been proposed by Cochet-Terrasson and Gaubert [19], Akian, Cochet-Terrasson, Detournay and Gaubert [20], and Bourque and Raghavan [21], allowing one to circumvent such degeneracies at the price of a complexification of the algorithm (handling the non-uniqueness of the bias). Hence, it is of interest to understand when such technicalities can be avoided.

Moreover, generally, the *existence* of the bias vector is controlled by some ergodicity conditions. A general result in [22] relates the existence of the bias vector of a Shapley operator with the uniqueness of the bias vector of an auxiliary (simpler) Shapley operator, called recession function. This provides further motivation to study the uniqueness problem for the bias vector.

C. Main results

We deal with two problems concerning zero-sum *parametric* stochastic games, that is, games whose parameters (transition payments and probabilities) vary.

We first look for conditions involving the transition probabilities which guarantee that the nonlinear eigenproblem is solvable *for all values of the transition payments*. Our first main result, Theorem 5, shows that this property is structural, meaning that it only depends on the support of the transition probabilities (not on their values).

Then, we address the question of the uniqueness of the bias vector, restricting our attention to games for which the previous structural property holds. Our second main result, Theorem 10, shows that the bias vector is generically unique (up to an additive constant). More precisely, we show that the set of transition payments for which the bias vector is not unique belongs to a polyhedral fan the cells of which have codimension one at least. A first ingredient in the proof is the max-plus spectral theorem, in the deterministic case, and its extension to the stochastic case [17]. A second ingredient is a general result, showing that the set of fixed points of a nonexpansive self-map of $\mathbb{R}^n$ is a retract of $\mathbb{R}^n$, see Theorem **??**. This allows us to infer the uniqueness of the bias vector of a Shapley operator from the uniqueness of the bias vector of the reduced Shapley operators obtained by fixing the strategy of one player.

Finally, as a simple application of our results, we show that the standard Hoffman and Karp policy iteration for two player games does terminate for generic well posed instances.

Let us note that although the finiteness of action spaces is assumed throughout the paper, the results in Section II remain true with compact action spaces (and some additional assumptions on the transition payments and probabilities, see [23]). In particular, they are true for repeated games with imperfect information (where players select an action simultaneously). However, the techniques used to prove the results in Section III are strongly related with the finiteness of action spaces. Therefore, they cannot go through Assumption A and the perfect information assumption.

The paper is organized as follows. In Section II, we show the structural existence result for the bias vector. This partly relies on recent results established by the authors in a companion work [23], that we first recall. The generic uniqueness of the bias vector is established in Section III, using results of max-plus spectral theory and nonlinear spectral theory. The application to policy iteration is presented in Section IV.

II. Structural properties for the existence of an eigenvalue

A. Necessary and sufficient condition of existence of an eigenvalue

Recall that, $\mathbb{R}^n$ being endowed with its usual partial order, the map $T : \mathbb{R}^n \to \mathbb{R}^n$ is said to be *monotone* and *additively homogeneous* if it satisfies the following properties, respectively,

- $x \leq y \;\Rightarrow\; T(x) \leq T(y),\;\; x, y \in \mathbb{R}^n$,
- $T(x + \lambda \mathbf{1}) = T(x) + \lambda \mathbf{1},\;\; x \in \mathbb{R}^n, \lambda \in \mathbb{R}$.

In [22], Gaubert and Gunawardena have given a sufficient condition for a monotone and additively homogeneous operator T on $\mathbb{R}^n$ to have an eigenvalue. This condition involves the *recession function* of T, denoted by $\hat{T}$ and defined on $\mathbb{R}^n$ by

$$\hat{T}(x) := \lim_{\alpha \to +\infty} \frac{T(\alpha x)}{\alpha}.$$

For a general operator, the recession function may not exist. However, if T is a Shapley operator (2) satisfying Assumption A (finite state and action spaces), then the recession function is well defined everywhere and given by

$$[\hat{T}(x)]_i = \min_{a \in A_i} \max_{b \in B_{i,a}} P_i^{ab} x, \;\; \forall i \in S. \tag{6}$$

Remark 1: If the recession function $\hat{T}$ of T exists, then it is also monotone and additively homogeneous. Furthermore, it is positively homogeneous, meaning that

- $\hat{T}(\lambda x) = \lambda \hat{T}(x),\;\; x \in \mathbb{R}^n, \lambda > 0$.

As a consequence, any vector proportional to $\mathbf{1}$ is a fixed point of $\hat{T}$.

Theorem 2 ([22]): Let $T : \mathbb{R}^n \to \mathbb{R}^n$ be a monotone and additively homogeneous operator. Suppose its recession function is well defined on $\mathbb{R}^n$ and has only trivial fixed points (that is, fixed point proportional to $\mathbf{1}$). Then T has an eigenvalue, meaning that Equation (5) is solvable for T.

Conversely, if F is a monotone, additively and positively homogeneous operator on $\mathbb{R}^n$ and if ν is one of its fixed points, then the operator $T : x \mapsto F(x) + \nu$ is monotone, additively homogeneous and satisfies $\hat{T} = F$ and $\chi(T) = \nu$. In particular, if ν is nontrivial, then T cannot have any eigenvalue. This leads to the following theorem.

Theorem 3: Let T be a Shapley operator (2) with finite state and action spaces (Assumption A). Then the following properties are equivalent:

i) T has an eigenvalue for all values of the transition payments r_i^{ab},
ii) the recession function of T (6) has only trivial fixed points.

Remark 4: The special case of Theorem 3 restrained to min-max functions (which corresponds to Shapley operators of deterministic games) has been addressed by Zhao, Zheng and Zhu [24]. In this case, the condition on the recession function can be reduced to a fixed point problem for a Boolean function (defined on the set $\{0,1\}^n$ of Boolean vectors), hence a purely combinatorial problem.

B. The condition of existence of an eigenvalue is structural

1) Statement of the result: In the following section, we shall consider operators (2) that have an eigenvalue for all values of the transition payments. According to Theorem 3, these are characterized by the property that the operator given by (6) has only trivial fixed points.

Fixed point problems related with such operators (6), which we call *payment-free* Shapley operators, have been studied in our companion work [23]. Our first main result is the following.

Theorem 5: Whether a given payment-free Shapley operator (6) has only trivial fixed points depends only on the support of the transition probabilities.

Recall that the support of a probability $P \in \Delta(S)$ is the subset defined as $\{j \in S \mid P_j \neq 0\}$. In other words, changing the probabilities P_i^{ab}, while leaving invariant the set of (i, j, a, b) such that $P_{ij}^{ab} \neq 0$ does not change the existence of a non-trivial fixed point of the operator $\hat{T}$ in (6).

We next recall some constructions and results that we use to establish Theorem 5.

2) Galois connection between invariant faces of $[0,1]^n$*:* Let $F : \mathbb{R}^n \to \mathbb{R}^n$ be a payment-free Shapley operator. If u is a fixed point of F then, denoting $I := \arg\min u$ and $J := \arg\max u$, we can show that

$$F(\mathbf{1}_{S\setminus I}) \leq \mathbf{1}_{S\setminus I} \quad \text{and} \quad \mathbf{1}_J \leq F(\mathbf{1}_J), \tag{7}$$

where $\mathbf{1}_K$, with $K \subset S$, is the vector with entries 1 on K and 0 on $S \setminus K$.

Let $\mathcal{F}^-$ (resp. $\mathcal{F}^+$) be the family of subsets of S verifying the first (resp. the second) inequality in (7):

$$\mathcal{F}^- := \big\{I \subset S \mid F(\mathbf{1}_{S\setminus I}) \leq \mathbf{1}_{S\setminus I}\big\},$$
$$\mathcal{F}^+ := \big\{J \subset S \mid \mathbf{1}_J \leq F(\mathbf{1}_J)\big\}.$$

The elements of $\mathcal{F}^-$ and $\mathcal{F}^+$ can be visualized as faces of the hypercube $[0,1]^n$ that are invariant by F.

The families $\mathcal{F}^-$ and $\mathcal{F}^+$ are nonempty lattices of subsets (they both contain $\emptyset$ and S). We define a Galois connection (see [25] for the definition) between them, denoted by $(\Phi, \Phi^\star)$, in the following way: to $I \in \mathcal{F}^-$ we associate the greatest element $J \in \mathcal{F}^+$ which has an empty intersection with I, and vice versa. Formally we have, for $I \in \mathcal{F}^-$ and $J \in \mathcal{F}^+$,

$$\Phi(I) := \bigcup_{J\in\mathcal{F}^+ : I\cap J=\emptyset} J \quad \text{and} \quad \Phi^\star(J) := \bigcup_{I\in\mathcal{F}^- : I\cap J=\emptyset} I. \tag{8}$$

If $I \in \mathcal{F}^-$, we denote by $\bar{I} := \Phi^\star \circ \Phi(I)$ its closure by the Galois connection (likewise, $\bar{J}$ is the closure of $J \in \mathcal{F}^+$).

The following results explain why this Galois connection is useful for the problem of existence of nontrivial fixed points.

Theorem 6 ([23]): Let F be a payment-free Shapley operator and $I \in \mathcal{F}^-$.

i) If $\Phi(I) = \emptyset$, then F has no nontrivial fixed point u such that $\arg\min u = I$.

ii) If $I \neq \emptyset$ and $I = \bar{I}$, then F has a fixed point u satisfying $\arg\min u = I$.

Corollary 7: A payment-free Shapley operator has a nontrivial fixed point if and only if there is a nontrivial subset $I \in \mathcal{F}^- \setminus \{\emptyset, S\}$ closed for the Galois connection $(\Phi, \Phi^\star)$, that is a set $I \in \mathcal{F}^- \setminus \{\emptyset, S\}$ such that $I = \bar{I}$.

3) Boolean abstractions of Shapley operators: Let F be a payment-free Shapley operator. We call *upper* and *lower Boolean abstractions* of F the operators defined on the set of Boolean vectors $\{0,1\}^n$, given respectively by

$$[F^+(x)]_i := \min_{a\in A_i} \max_{b\in B_{i,a}} \max_{j:P_{ij}^{ab}>0} x_j,$$
$$[F^-(x)]_i := \min_{a\in A_i} \max_{b\in B_{i,a}} \min_{j:P_{ij}^{ab}>0} x_j,$$

for every $i \in S$. Note that these operators depend only on the support of the transition probabilities.

They suffice to characterize the families $\mathcal{F}^-$ and $\mathcal{F}^+$, as shown by the following lemma.

Lemma 8 ([23]): Let F be a payment-free Shapley operator and let $I, J \subset S$. Thenn

$$I \in \mathcal{F}^- \;\Leftrightarrow\; F^+(\mathbf{1}_{S\setminus I}) \leq \mathbf{1}_{S\setminus I},$$
$$J \in \mathcal{F}^+ \;\Leftrightarrow\; F^-(\mathbf{1}_J) \geq \mathbf{1}_J.$$

Corollary 9: Given a payment-free Shapley operator, the families $\mathcal{F}^-$ and $\mathcal{F}^+$ and the Galois connection $(\Phi, \Phi^\star)$ are uniquely determined by the supports of the transition probabilities.

III. Generic uniqueness of the bias vector of zero-sum games

A. The setting

We fix the state space, the action spaces and the transition probabilities and consider parametric Shapley operators (2), that is, operators whose transition payments (thought of as parameters) can vary. In particular, we denote by $\mathcal{R}$ the *space of transition payments*, whose elements are the finite collections of reals $\big(r_i^{ab}\big)_{(i,a,b)}$ for the $(i,a,b) \in S \times A \times B$ such that $a \in A_i$ and $b \in B_{i,a}$. This set can be identified with the Euclidean space $\mathbb{R}^q$ for a certain integer q. For a given $r \in \mathcal{R}$, we denote by T_r the corresponding Shapley operator as defined in (2).

In this section we assume that the eigenproblem (5) is solvable for all the values $r \in \mathcal{R}$ of the transition payments. For any Shapley operator T as defined in (2), we denote by $\mathcal{E}(T)$ the set of its eigenvectors, and we say that T has a unique eigenvector (up to an additive constant) if $\mathcal{E}(T)$ is reduced to a line (which is invariant by the addition of a constant).

Before stating the main result of this section, we recall that a polyhedral fan in $\mathbb{R}^q$ is a polyhedral complex consisting of polyhedral cones, that is, a finite set $\mathcal{K}$ of polyhedral cones such that

- $P \in \mathcal{K}$ and F is a face of P implies that $F \in \mathcal{K}$,
- for all $P, Q \in \mathcal{K}$, $P \cap Q$ is a face of P and Q.

We refer to the textbook [26] for background on polyhedral complexes and fans (but be aware that in this book a polyhedral complex is not necessarily finite).

Theorem 10: Assume that for any transition payment $r \in \mathcal{R}$, the Shapley operator T_r has an eigenvalue. Then the space $\mathcal{R}$ is covered by a polyhedral fan such that for each r in the interior of a full-dimensional cone, T_r has a unique eigenvector.

Note that Theorem 10 implies that the transition payments $r \in \mathcal{R}$ for which T_r has more that one eigenvector (up to an additive constant) are situated in the finite union of subspaces of codimension 1.

We prove Theorem 10 in the framework of both deterministic and stochastic games. The structure of the proofs are similar but they rely on different theories, which we find interesting to present, although deterministic games can be seen as a particular case of stochastic games. These theories are max-plus algebra and nonlinear Perron-Frobenius theory, respectively.

However, in both cases we infer from the following result a fundamental property about the geometry of the set of eigenvectors.

Theorem 11 ([27]): Let C be a closed convex subset of a finite dimensional Banach space. Let $T : C \to C$ be a nonexpansive map with respect to the norm of the Banach space. Then the fixed point set of T is a nonexpansive retract of C, meaning that it is the image of a nonexpansive projection of C.

Corollary 12: Let $T : \mathbb{R}^n \to \mathbb{R}^n$ be an monotone and additively homogeneous map (hence nonexpansive in the sup-norm). Suppose Equation (5) is solvable. Then $\mathcal{E}(T)$ is a retract of $\mathbb{R}^n$ by a nonexpansive map. In particular, $\mathcal{E}(T)$ is arcwise connected.

B. Deterministic case

In the deterministic case, each transition probability P_i^{ab} has only one state in its support. Thus, for $r \in \mathcal{R}$, the operator T_r in (2) can be written as

$$[T_r(x)]_i = \min_{a\in A_i} \max_{1\le j\le n} \left(\tilde{r}_{ij}^a + x_j\right), \tag{9}$$

where $\tilde{r}_{ij}^a = \max\{r_i^{ab} \mid b \in B_{i,a},\ P_{ij}^{ab} = 1\}$ with the convention that $\max \emptyset = -\infty$. Moreover, the set $C_{i,a} := \{j \in S \mid \tilde{r}_{ij}^a \in \mathbb{R}\}$ depends only on the transition probabilities, hence is independent of the value of the payments.

The above representation (9) can be simplified using max-plus notations. For this purpose, we introduce the max-plus semiring $\mathbb{R}_{\max} := \mathbb{R} \cup \{-\infty\}$, which is a commutative idempotent semiring with addition $x \oplus y = \max(x, y)$ and multiplication $x \otimes y = x + y$. The zero and unit elements of $\mathbb{R}_{\max}$ are $-\infty$ and 0, respectively.

Let Σ be the (finite) set of policies of player MIN, that is, the maps $\sigma : S \to A, i \mapsto \sigma(i) \in A_i$. If $\mathbb{R}^n$ is endowed with its usual partial order then, for $r \in \mathcal{R}$, the operator T_r can be written as

$$T_r(x) = \min_{\sigma\in\Sigma} M_r^\sigma x, \tag{10}$$

where $M_r^\sigma \in \mathbb{R}_{\max}^{n\times n}$ is a max-plus matrix (that is a $n \times n$ matrix over the max-plus semiring $\mathbb{R}_{\max}$) and $M_r^\sigma x$ is to understand in the max-plus sense. Precisely, we have $[M_r^\sigma]_{ij} = \tilde{r}_{ij}^{\sigma(i)}$. Note that (10) means that, for every $x \in \mathbb{R}^n$, we have $T_r(x) \le M_r^\sigma x$ for all $\sigma \in \Sigma$, and there is $\sigma \in \Sigma$ such that $T_r(x) = M_r^\sigma x$. In particular, if λ is the eigenvalue of T_r and u is an eigenvector, then there is a policy σ such that $M_r^\sigma u = \lambda\mathbf{1} + u$.

1) Max-plus algebra: We now present some results on max-plus matrices that can be found in [9, Chapter 3] or [13]. We draw the attention to the fact that vectors in $\mathbb{R}_{\max}^n$ are allowed to have entries equal to $-\infty$.

Let $M \in \mathbb{R}_{\max}^{n\times n}$ with no row identically equal to $-\infty$. We need the following definitions.

Definition 13: The *precedence graph* of M, denoted by $\mathcal{G}(M)$, is the weighted directed graph with set of nodes S and an arc from i to j if $M_{ij} \neq -\infty$, in which case the weight of the arc is the real number M_{ij}.

Definition 14: The *maximal circuit mean*, $\rho(M)$, is the maximum average weight of the elementary circuits of $\mathcal{G}(M)$ (the circuits where every nodes are distinct, except the starting and the ending one).

Theorem 15 ([9, Rem. 3.24]): Let $M \in \mathbb{R}_{\max}^{n\times n}$ with no row identically equal to $-\infty$. Then, there exists $u \in \mathbb{R}_{\max}^n$, with at least one finite entry, such that $Mu = \rho(M)\mathbf{1} + u$.

The vocabulary is the same as for the eigenproblem (5): we say that u is a tropical eigenvector of M associated with the tropical eigenvalue $\rho(M)$.

Definition 16: The *critical graph* of M, denoted by $\mathcal{G}^c(M)$, is the subgraph of $\mathcal{G}(M)$ consisting of all the elementary circuits with average weight equal to $\rho(M)$. A critical class of M is a connected component of $\mathcal{G}^c(M)$.

Theorem 17 ([9, Thm. 3.101]): Let $M \in \mathbb{R}_{\max}^{n\times n}$ with no row identically equal to $-\infty$. Denote by $\mathcal{E}_{\max}(M)$ the set of tropical eigenvectors of M associated with the eigenvalue $\rho(M)$. Then, $\mathcal{E}_{\max}(M)$ is tropically generated by a finite family of vectors in $\mathbb{R}_{\max}^n$. Moreover, the least cardinality of a generating family of $\mathcal{E}_{\max}(M)$ is equal to the number of critical classes of M.

Theorem 18: Let $M \in \mathbb{R}_{\max}^{n\times n}$ and $\lambda \in \mathbb{R}$. If there is a vector $u \in \mathbb{R}^n$ such that $Mu = \lambda\mathbf{1} + u$, then $\lambda = \rho(M)$.

2) Sketch of proof of Theorem 10: By definition of $\mathcal{R}$, for every σ, the matrices M_r^σ have the same precedence graph for all transition payments $r \in \mathcal{R}$: there is an arc from i to j if $j \in C_{i,\sigma(i)}$. Then, for every σ, we can define the piecewise linear function that maps an element $r \in \mathcal{R}$ to the maximal circuit mean of M_r^σ, $\rho(M_r r^\sigma)$, and consider the polyhedral fan $\mathcal{C}^\sigma$ associated with the cone of linearity of this function. To each cone of linearity with full dimension corresponds a unique critical circuit in $\mathcal{G}(M_r^\sigma)$ (where r is any transition payment).

We consider then the polyhedral fan $\mathcal{C}$ obtained as the intersection of all the fans $\mathcal{C}^\sigma$. We show that for each element $r \in \mathcal{R}$ in the interior of a full-dimensional cone of $\mathcal{C}$, an eigenvector of T_r is also a tropical eigenvector of a certain M_r^σ, associated with the eigenvalue $\rho(M_r^\sigma)$, and whose tropical eigenvector is unique (up to an additive constant). This shows that the eigenspace $\mathcal{E}(T_r)$ contains a finite number of elements up to the addition of constants. The conclusion follows from Theorem 12.

C. Stochastic case

We go back to the general framework presented in Section III-A. Recall that Σ is the set of policies of player MIN. For all $\sigma \in \Sigma$, we define Π^σ as the set of policies of player MAX, when the policy of MIN is fixed to σ, that is the maps $\pi : S \to B, i \mapsto \pi(i) \in B_{i,\sigma(i)}$. Moreover, given $\sigma \in \Sigma$ and $r \in \mathcal{R}$, we define the operator $T_r^\sigma : \mathbb{R}^n \to \mathbb{R}^n$ by

$$[T_r^\sigma(x)]_i = \max_{b\in B_{i,\sigma(i)}} \left(r_i^{\sigma(i)b} + P_i^{\sigma(i)b} x\right),$$

for all $i \in S$. Then, we have ($\mathbb{R}^n$ being endowed with its usual partial order)

$$T_r(x) = \min_{\sigma\in\Sigma} T_r^\sigma(x). \tag{11}$$

Likewise, given $\sigma \in \Sigma$, we have

$$T_r^\sigma(x) = \max_{\pi\in\Pi^\sigma} \left(r^{\sigma\pi} + P^{\sigma\pi} x\right), \tag{12}$$

where $r^{\sigma\pi}$ is the vector in $\mathbb{R}^n$ whose i-th entry is defined by $r_i^{\sigma\pi} = r_i^{\sigma(i)\pi(i)}$ and $P^{\sigma\pi}$ is the $n \times n$ stochastic matrix whose i-th row is given by $P_i^{\sigma\pi} = P_i^{\sigma(i)\pi(i)}$.

To prove Theorem 10, we need to state some results on operators of type (12), which are piecewise affine, convex (componentwise), monotone and additively homogeneous.

1) Ergodic stochastic control: Firstly, we characterize the eigenvalue of a componentwise convex operator T^σ as defined in (12).

Lemma 19: Let T^σ be an operator of type (12). For every $\pi \in \Pi^\sigma$, let $m_C^{\sigma\pi}$ be the invariant measure associated with the final class C of the stochastic matrix $P^{\sigma\pi}$. Denote by $\mathcal{F}^{\sigma\pi}$ the set of final classes of $P^{\sigma\pi}$. Suppose that T^σ has an eigenvalue and denote it by λ. Then,

$$\lambda = \max_{\pi \in \Pi^\sigma,\ C \in \mathcal{F}^{\sigma\pi}} \langle m_C^{\sigma\pi}, r^{\sigma\pi} \rangle.$$

Secondly, we characterize the set of eigenvectors of the operator T^σ. For that, we need few definitions that are detailed in [17]. For a map $F : \mathbb{R}^n \to \mathbb{R}^n$, componentwise convex, the notion of subdifferential is generalized by setting, for $x \in \mathbb{R}^n$, $\partial F(x) = \{P \in \mathbb{R}^{n\times n} \mid P(y-x) \leq F(y) - F(x),\ \forall y \in \mathbb{R}^n\}$. If F is monotone and additively homogeneous, it can be shown that $\partial F(x)$ is a convex set of stochastic matrices. Suppose that F has an eigenvector u. Then the critical graph of F, denoted by $\mathcal{G}^c(F)$, is the directed graph defined as the union of the final graphs of the stochastic matrices $P \in \partial F(u)$. This graph does not depend on the choice of the eigenvector u. A critical node of F is a node of $\mathcal{G}^c(F)$ and a critical class of F is the set of nodes of a strongly connected component of $\mathcal{G}^c(F)$.

Theorem 20 ([17, Thm 1.1]): Let $F : \mathbb{R}^n \to \mathbb{R}^n$ be a convex, monotone and additively homogeneous map. Suppose that F has an eigenvector and denote by C the set of its critical nodes. Then, (i) the restriction $\mathsf{r} : \mathbb{R}^n \to \mathbb{R}^C, x \mapsto (x_i)_{i\in C}$ is an isomorphism from $\mathcal{E}(F)$ to its image $\mathcal{E}^c(F)$; (ii) $\mathcal{E}^c(F)$ is a convex set whose dimension is at most equal to the number of critical classes of F, and this bound is attained when F is piecewise affine.

2) Sketch of proof of Theorem 10: Let $r \in \mathcal{R}$. For every σ, the operaor T_r^σ is convex, monotone and additively homogeneous. Then, for every σ, we can define the piecewise linear map $\rho^\sigma : r \in \mathcal{R} \mapsto \max_{\pi\in\Pi^\sigma,\ C\in\mathcal{F}^{\sigma\pi}} \langle m_C^{\sigma\pi}, r^{\sigma\pi}\rangle$, where, for $\pi \in \Pi^\sigma$, $m_C^{\sigma\pi}$ is the invariant measure associated with the final class C of the stochastic matrix $P^{\sigma\pi}$. We consider the polyhedral fan $\mathcal{C}^\sigma$ associated with the cone of linearity of the map ρ^σ. To each cone of linearity with full dimension corresponds a unique final class C.

We consider then the polyhedral fan $\mathcal{C}$ obtained as the smallest refinement of all the fans $\mathcal{C}^\sigma$. We show that for each transition payment r in the interior of a full-dimensional cone of $\mathcal{C}$, an eigenvector u of T_r is also an eigenvector of a certain T_r^σ with a unique critical class, hence a unique eigenvector (up to an additive constant). This shows that the eigenspace $\mathcal{E}(T_r)$ contains a finite number of elements up to the addition of constants. The conclusion follows from Theorem 12.

IV. Application to policy iteration

Let us use the notations of Section III-C. In particular, Σ is the (finite) set of policies of player MIN, that is the maps $\sigma : S \to A, i \mapsto \sigma(i) \in A_i$, Π^σ is the (finite) set of policies of player MAX when the policy of MIN, $\sigma \in \Sigma$, is fixed, that is the set of maps $\pi : S \to B, i \mapsto \pi(i) \in B_{i,\sigma(i)}$, and for all $\sigma \in \Sigma$, $\pi \in \Pi^\sigma$, $P^{\sigma\pi}$ is the $n \times n$ stochastic matrix whose ith row is given by $P_i^{\sigma\pi} = P_i^{\sigma(i)\pi(i)}$.

When T is a Shapley operator (2) with finite state and action spaces (Assumption A) such that all the stochastic matrices $P^{\sigma\tau}$ are irreducible, Hoffman and Karp [18] introduced the following algorithm with input T and output an eigenvalue λ and a bias vector u of T. Then optimal stationary policies for both players can be computed.

Algorithm 21 ([18]): Select an arbitrary policy $\sigma_0 \in \Sigma$ of player MIN, then apply successively the two following steps for $k \geq 0$, until $\sigma_{k+1} = \sigma_k$:

i) Compute the eigenvalue λ^k and the bias v^k of the game with fixed policy σ_k of player MIN, that is the solutions λ and v of $\lambda\mathbf{1} + v = T^{\sigma_k}(v)$;
ii) Improve the policy: choose an optimal policy for v^k, that is $\sigma_{k+1} \in \Sigma$ such that $T(v^k) = T^{\sigma_{k+1}}(v^k)$, with $\sigma_{k+1} = \sigma_k$ as soon as this is possible.

Return λ^k and v^k.

In the above algorithm, Step i) is solved by using the same policy iteration for the (one-player) game with fixed policy σ_k, which constructs $\lambda^{k,\ell}$, $v^{k,\ell}$ and $\pi_{k,\ell}$ from some $\pi_{k,0}$. One can show that the sequence λ^k of Algorithm 21 is nonincreasing ($\lambda^{k+1} \leq \lambda^k$), and that when $\lambda^{k+1} = \lambda^k$, then $v^{k+1} = v^k$ (up to an additive constant) and $\sigma_{k+1} = \sigma_k$. Hence, the algorithm terminates after a finite number of steps, for the set of policies Σ is finite.

When the stochastic matrices $P^{\sigma\tau}$ are not irreducible, the existence of λ^k and v^k may fails in Step i) of Algorithm 21. Moreover, even if λ^k and v^k can be constructed at each step k, the nonuniqueness of v^k (up to an additive constant) may lead to a cycling of the sequence v^k in Algorithm 21, as shown for instance in [20, Section 6]. In the general case, one need first to replace eigenvectors by invariant half-lines to avoid problems of existence and second to apply a special treatment of degeneracies (when the eigenvalue or mean payoff vector does not change from one step to another and the bias vector or invariant half-line is not unique) as in [19], [20], [21] to avoid cycling of the algorithm.

Using the arguments of Section III-C, we obtain the following result, under the same setting as for Theorem 10, that is fixing the state space, the action spaces and the transition probabilities and considering parametric Shapley operators (2).

Theorem 22: The space $\mathcal{R}$ is covered by a polyhedral fan such that for each transition payment r in the interior of a full-dimensional cone, if Algorithm 21 is well posed, that is if at each step k of the algorithm, the operator $T_r^{\sigma_k}$ has an eigenvalue, then it terminates after a finite number of steps and gives an eigenvalue λ and a bias vector u of T_r.

Sketch of proof: Considering the same fan $\mathcal{C}$ as in Section III-C, we get that for each r in the interior of a full-dimensional cone of $\mathcal{C}$, and for each $\sigma \in \Sigma$, if the operator T_r^σ has an eigenvector u, then this eigenvector is unique (up to an additive constant). This implies that, if at each step k of the algorithm, the operator $T_r^{\sigma_k}$ has an eigenvalue λ^k, then the bias vector v^k is unique up to an additive constant. One can always show (without any assumption) that the sequence λ^k is nonincreasing ($\lambda^{k+1} \leq \lambda^k$). Then, using the uniqueness of bias vectors, we deduce from Lemma 3.3 of [17], that if $\lambda^{k+1} = \lambda^k$, then up to an additive constant, $v^{k+1} \leq v^k$ with equality on the set of critical nodes of $T_r^{\sigma_k}$. This implies that the sequence v^k coincides up to an additive constant with the sequence obtained in the algorithm introduced in [19] and developed in [20]. Then, the convergence of Algorithm 21 in a finite number of steps follows from the convergence of the algorithm of [19], [20].

V. Example

Consider the following parametric Shapley operator defined on $\mathbb{R}^3$ (here we use $\wedge$ and $\vee$ instead of $\min$ and $\max$, respectively, and we recall that $+$ has precedence over them):

$$T_r(x) = \begin{pmatrix} r_1 + \frac{1}{2}(x_1+x_3) \wedge r_2 + \frac{1}{2}(x_1+x_2) \\ r_3 + \frac{1}{2}(x_1+x_3) \wedge \left(r_4 + \frac{1}{2}(x_1+x_2) \vee r_5 + x_3\right) \\ r_6 + \frac{1}{2}(x_1+x_3) \vee r_7 + x_3 \end{pmatrix}$$

with $r \in \mathbb{R}^7$. Its recession function is

$$\hat{T}_r(x) = \begin{pmatrix} \frac{1}{2}(x_1+x_3) \wedge \frac{1}{2}(x_1+x_2) \\ \frac{1}{2}(x_1+x_3) \wedge \left(\frac{1}{2}(x_1+x_2) \vee x_3\right) \\ \frac{1}{2}(x_1+x_3) \vee x_3 \end{pmatrix}$$

and it can be checked that it has only trivial fixed points. Hence, according to Theorem 3, Equation (5) is solvable for every $r \in \mathbb{R}^7$. Alternatively, we may consider the Galois connection between the families $\mathcal{F}^- = \{\emptyset, \{1,3\}, \{1,2,3\}\}$ and $\mathcal{F}^+ = \{\emptyset, \{3\}, \{1,2,3\}\}$ as defined in II-B.2. We verify that the only nontrivial element of $\mathcal{F}^-$ is not closed with respect to this Galois connection (since $\Phi^\star \circ \Phi(\{1,3\}) = \Phi^\star(\emptyset) = \{1,2,3\}$), which leads to the same conclusion according to Corollary 7.

We now fix $r = r_0 = (0,1,2,1,-2,-3,1)$ and denote by T_0 the corresponding operator. We consider the perturbations of r_0 that depend only on the state i.e., in the space of payments $\mathbb{R}^7$, we consider the affine subspace $\mathcal{A}_0 = \{r_0 + (g_1, g_1, g_2, g_2, g_2, g_3, g_3) \mid g \in \mathbb{R}^3\}$, corresponding to the operators $T_0 + g$ with $g \in \mathbb{R}^3$ (note that by additive homogeneity of T_r, we may assume w.l.o.g. that $g_3 = 0$). Figure 1 shows the polyhedral complex obtained as the intersection between the polyhedral fan of Theorem 10 and the affine subspace $\mathcal{A}_0$. Here for each g in the interior of a full-dimensional polyhedron, T_0+g has a unique eigenvector (up to an additive constant). Let us detail what happens in the neighborhood of $g = 0$, point in which $T_0 + g$ fails to have a unique eigenvector. Note that in the neighborhood of $g = 0$, the eigenvalue of $T_0 + g$ remains 1.

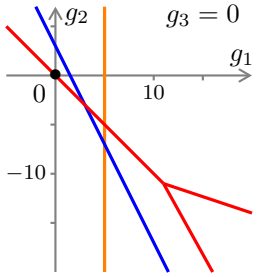

Fig. 1.

If $g_1 + g_2 = 0$, the eigenvectors of $T_0 + g$ are defined by $x_1 = x_2 + 2g_1$ and $-3 + g_2 \leq x_2 - x_3 \leq -2 - g_1$.
If $g_1 + g_2 > 0$, the unique eigenvector (up to an additive constant) is $(-2 + 2g_1, -2 + 2g_1 + 2g_2, 0)$.
If $g_1 + g_2 < 0$, the unique eigenvector is $(-3 + 2g_1 + g_2, -3 + g_2, 0)$.

References

[1] A. Neyman and S. Sorin, *Stochastic Games and Applications*, ser. Nato Science Series C. Kluwer Academic Publishers, 2003, vol. 570.
[2] E. Kohlberg, "Invariant half-lines of nonexpansive piecewise-linear transformations," *Math. Oper. Res.*, vol. 5, no. 3, pp. 366–372, 1980.
[3] A. Neyman, "Stochastic games and nonexpansive maps," in *Stochastic Games and Applications*, ser. Nato Science Series C, A. Neyman and S. Sorin, Eds. Kluwer Academic Publishers, 2003, vol. 570, ch. 26, pp. 397–415.
[4] D. Rosenberg and S. Sorin, "An operator approach to zero-sum repeated games," *Israel J. Math.*, no. 121, pp. 221–246, 2001.
[5] I. V. Romanovski, "Optimization of stationary control of discrete deterministic process in dynamic programming," *Kibernetika*, vol. 3, no. 2, pp. 66–78, 1967.
[6] M. Gondran and M. Minoux, "Valeurs propres et vecteurs propres dans les dioïdes et leur interprétation en théorie des graphes," *Bull. Dir. Et. Rech. EDF, Série C*, pp. 25–41, 1977.
[7] R. Cuninghame-Green, *Minmax Algebra*. Springer-Verlag, Berlin, 1979.
[8] V. P. Maslov and S. N. Samborski, *Idempotent Analysis*, ser. Advances In Soviet Mathematics. Amer. Math. Soc., Providence, 1992, vol. 13.
[9] F. Baccelli, G. Cohen, G. Olsder, and J.-P. Quadrat, *Synchronization and Linearity: An Algebra for Discrete Event Systems*, ser. Wiley Series in Probability and Mathematical Statistics. John Wiley, 1992.
[10] R. B. Bapat, "A max version of the Perron-Frobenius theorem," *Linear Algebra Appl.*, vol. 275/276, pp. 3–18, 1998.
[11] L. A. Kontorer and S. Y. Yakovenko, "Nonlinear semigroups and infinite horizon optimization," in *Idempotent Analysis*, ser. Advances In Soviet Mathematics. Amer. Math. Soc., Providence, 1992, vol. 13, pp. 167–210.
[12] V. N. Kolokoltsov and V. P. Maslov, *Idempotent Analysis and Applications*. Kluwer Academic Publishers, 1997.
[13] M. Akian, R. B. Bapat, and S. Gaubert, "Max-plus algebras," in *Handbook of Linear Algebra, Second Edition*, ser. Discrete Mathematics and its Applications, L. Hogben, Ed. Chapman & Hall/CRC, Boca Raton, FL, 2013, ch. 35.
[14] B. Sturmfels and N. M. Tran, "Combinatorial types of tropical eigenvectors," *Bull. Lond. Math. Soc.*, vol. 45, no. 1, pp. 27–36, 2013.
[15] M. Akian, S. Gaubert, and C. Walsh, "The max-plus Martin boundary," *Documenta Mathematica*, vol. 14, pp. 195–240, 2009.
[16] A. Fathi, "Weak Kam theorem in Lagrangian dynamics," to appear.
[17] M. Akian and S. Gaubert, "Spectral theorem for convex monotone homogeneous maps and ergodic control," *Nonlinear Analysis, T.M.A.*, vol. 52, no. 2, pp. 637–679, 2003.
[18] A. J. Hoffman and R. M. Karp, "On nonterminating stochastic games," *Management Sci.*, vol. 12, pp. 359–370, 1966.
[19] J. Cochet-Terrasson and S. Gaubert, "A policy iteration algorithm for zero-sum stochastic games with mean payoff," *C. R. Math. Acad. Sci. Paris*, vol. 343, no. 5, pp. 377–382, 2006.
[20] M. Akian, J. Cochet-Terrasson, S. Detournay, and S. Gaubert, "Policy iteration algorithm for zero-sum multichain stochastic games with mean payoff and perfect information," arXiv:1208.0446.
[21] M. Bourque and T. Raghavan, "Policy improvement for perfect information and additive reward additive transition stochastic games with discounted and average payoff," 2012, preprint available from the author's web page, University of Illinois at Chicago.
[22] S. Gaubert and J. Gunawardena, "The Perron-Frobenius theorem for homogeneous, monotone functions," *Trans. Amer. Math. Soc.*, vol. 356, pp. 4931–4950, 2004.

[23] M. Akian, S. Gaubert, and A. Hochart, "Ergodicity conditions for zero-sum games," 2014, arXiv:1405.4658.
[24] Q. Zhao, D.-Z. Zheng, and X. Zhu, "Structure properties of min-max systems and existence of global cycle time," *IEEE Trans. Automat. Contr*, vol. 46, no. 1, pp. 148–151, 2001.
[25] G. Birkhoff, *Lattice Theory*, ser. Colloquium publications. American Mathematical Society, Providence, 1995 (first edition, 1940), vol. 25.
[26] J. D. Loera, J. Rambau, and F. Santos, *Triangulations, Stuctures for Algorithms and Applications*. Springer, 2010.
[27] R. Bruck, Jr., "Properties of fixed-point sets of nonexpansive mappings in Banach spaces," *Trans. Amer. Math. Soc.*, vol. 179, no. 251-262, 1973.